

5/2-32
146842
P-17
February 15, 1993

N93-24671

Management of Space Networks

R. W. Markley and B. F. Williams
Advanced Information Systems Section

NASA has proposed missions to the Moon and Mars that reflect three areas of emphasis: human presence, exploration, and space resource development for the benefit of Earth. A major requirement for such missions is a robust and reliable communications architecture. Network management—the ability to maintain some degree of human and automatic control over the span of the network from the space elements to the end users on Earth—is required to realize such robust and reliable communications. This article addresses several of the architectural issues associated with space network management.

Round-trip delays, such as the 5- to 40-min delays in the Mars case, introduce a host of problems that must be solved by delegating significant control authority to remote nodes. Therefore, management hierarchy is one of the important architectural issues.

The following article addresses these concerns, and proposes a network management approach based on emerging standards that covers the needs for fault, configuration, and performance management, delegated control authority, and hierarchical reporting of events. A relatively simple approach based on standards was demonstrated in the DSN 2000 Information Systems Laboratory, and the results are described.

I. Introduction

NASA has proposed missions to the Moon and Mars that reflect three areas of emphasis: human presence, exploration, and space resource development for the benefit of Earth. The Moon is a natural test-bed to prepare for missions to Mars through simulation, systems testing, operations, and studying human capabilities.

Communications with the Moon should be relatively straightforward with existing Deep Space Network (DSN) systems because lunar operations are initially planned to be centered at a main base on the near side of the Moon.

However, a notable design issue will be communications from remote sites on the Moon. Such issues will increase the complexity of the space network.

The report on America's Space Exploration Initiative, *America at the Threshold* [1] describes the increased complexity of Martian communications as follows:

"Providing communications for the Martian missions is considerably more challenging than for lunar missions. Mars can be as much as 1,000 times more distant from Earth than the Moon, which results in a spatial signal loss one million times greater. In addition, Mars rotates at

about the same rate as the Earth, putting surface locations out of direct touch for over 12 hours at a time."

A tentative communications architecture for Mars (Fig. 1) proposes Mars-synchronous relay satellites to provide continuous coverage of surface elements as well as orbital elements [1]. The Martian main base may also have the ability to communicate directly with Earth when in view. Again in this case, the complexity of the space network will be increased as mobility and distribution of space elements are factored into the mission design.

A major requirement for such a lunar and Martian communications system is network management—the need to maintain some degree of human and automatic control over the network to assure highly reliable and robust communications from the space elements to the end users.

The following article addresses these issues, and proposes a network management architecture based on emerging commercial technology that covers the needs for fault, configuration, and performance management, delegated control authority, and hierarchical reporting of events. An approach based on standards was demonstrated in the DSN 2000 Information Systems Laboratory, and the results are described in this article.

II. Requirements

A. Space Management Network

Figure 2 illustrates a simplified schematic of the primary network to the Moon and Mars that will support science and human exploration. The figure illustrates the multiple paths between data sources and destinations, and the redundancy that is built into the primary architecture. For example, communications between one of the DSN antennas at a Deep Space Communications Complex (DSCC) and the Mars main base (Fig. 2(b)) may be direct or routed through a relay satellite. Data from the Mars habitat to a remote scientific instrument on Mars may be direct or routed through facilities on the relay satellite.

Round-trip delays, such as the 5- to 40-minute delays in the Mars case, introduce a host of problems that must be solved by delegating significant control authority to remote nodes.

This article proposes a Space Management Network (SMN) to support transmission of *management* data from all the network elements (from the space elements to the end users) to Earth-based operations centers. The SMN is a logical network that can be distinguished from the primary network because its chief function is to support

transmission of management data. The SMN may use dedicated facilities or share facilities with the primary network.

Major nodes in the SMN include the mission and science support centers, the DSN, the lunar main base, and the Mars main base. The SMN may interface to each of these processing end points, as well as intermediate communication facilities.

In some ways space network management is similar to the management of complex Earth-based communications networks in which there are many types of interconnected networks, such as local area networks (LANs) and wide area networks (WANs). (In our unique network, however, there are also space segments.) Management of these complex configurations is an area of current research and development because of (1) the large numbers of nodes (in the thousands for some enterprises), (2) the geographic distances involved, (3) the remoteness of much of the equipment, (4) the need for human management of the systems, and (5) the vision that it is now an achievable goal because of the recent standardization of network management protocols.

B. Domains

The extent of the primary space network suggests a logical separation into four domains: Earth, Space, Moon, and Mars. Each domain may require varying degrees of security, performance, and availability. The networks in each domain are summarized in Table 1. This table identifies the types of subnets; the actual numbers of such subnets are a detailed design issue that will evolve with further mission planning.

1. Earth Domain. The Earth domain includes all the networks uniting the the network, mission, and science operations centers. The Deep Space Network, as the likely focus of network operations, would route video, voice, and data to the mission and science operations centers over domestic and international circuits. The network operations center could also be the primary location for integrated management of the end-to-end network.

Mission operations and science operations centers are responsible for management of the mission and monitoring and analyzing scientific data. A reasonable assumption is that thousands of network elements would be included in the Earth domain.

2. Space Domain. Moon and Mars main bases will be the destination end of the direct link from Earth. The main bases are likely to be the primary human interface for local network management activities.

Remote Moon rovers and science instruments may communicate directly with Earth if they are beyond the line of sight to the main base.

A second mode of Mars-Earth communication will be Mars-relay-satellite to Earth. This mode will be used when the main base is unable to communicate to the Earth because of Mars' planetary rotation. These links will require a complex management system with redundant equipment and redundant command channels.

It has been suggested that a solar activity warning system may be necessary to protect human explorers.¹ Such emergency data would be reported directly to the Moon and Mars over dedicated links. Depending on the Mars permanent outpost position, the data may go directly to the outpost or be routed through the communications relay. The data would provide timely warning of solar activity that may pose health hazards to humans on the surface. Such solar bursts induce extreme noise on the links and the links themselves will require substantial fault and error protection.

3. Moon Domain. Lunar surface-to-surface communications would be used for video, voice, and data communications. Nodes include a habitat, remote science instruments, and mobile rovers and humans involved in extravehicular activity (EVA). Locally, the links may be wire, optical fiber, or line-of-sight radio. Fault and configuration management are major issues with lunar communications.

Beyond the line of sight (approximately 6 km with a 10-m high antenna), lunar radio communications require a lunar surface path with intermediate radio relays, or a round-about path via direct Earth links.

4. Mars Domain. Surface-to-surface communications would be used for video, voice, and data communications. Nodes are similar to lunar nodes with similar communications options. Fault and configuration management are also major issues with Martian communications.

Beyond the line of sight (approximately 8 km with a 10-m high antenna), communications would likely be over intermediate surface radio relays or through an orbiting satellite relay. Surface-to-relay satellite communications would provide a path not only to locations over the horizon, but to the Earth and to other orbiting spacecraft.

Relay-satellite-to-relay-satellite communications extend the routing of voice and data to Earth beyond the

limits of just one relay satellite. These links would also require a complex network management system with redundant equipment and redundant command channels to Mars and Earth. Manned and unmanned orbiting spacecraft at Mars will have network requirements similar to the communication relays. Manned orbiters will require intensive communication to operations centers on Earth and local communications to Mars.

III. Network Management Technology

It is the premise of this article that end-to-end network management will be accomplished through a structured, evolvable management architecture based on standards because such an architecture is likely to minimize life-cycle costs. Hopefully, if the proper standards are chosen, a standards-based architecture will lead to utilization of low-cost commercial software products. The following elements are essential for the description of this architecture: (1) management model, (2) protocol architecture, (3) connectivity, and (4) human interfaces.

A. Management Model

In concept, network management usually involves application processes called "managers" on managing systems and "agents" on managed systems. Current commercial approaches generally focus on a management hierarchy with three layers of control (Fig. 3): element managers, network managers, and an integrated network manager.

An *element manager* performs management functions relative to that communication element and displays the data locally or makes the data available to a higher level management system. Typically network elements are managed through a *software agent*. The agent is devoted to monitoring the status and activity of the network element. The agent may be periodically polled by a higher level system or initiate urgent messages to a designated system if some threshold has been exceeded. If the concept were applied to the primary space network, examples would include agent software in network elements of the lunar main base or the DSN ground systems.

A *network manager* usually manages multiple elements on one type of network. A network manager is also concerned with managing the network circuits or channels, monitoring such parameters as link utilization and total packet rate over the medium. The area of surveillance could range in size from an area as small as a science data processing laboratory to as large as an antenna complex.

An *integrated network manager* is a layer of control at the highest layer that can incorporate information from

¹ W. Kurth, personal communication, University of Iowa, January 5, 1992.

many types of networks. It typically interfaces to many network managers. An integrated network manager could be used to oversee end-to-end communications on the Earth, in space, and on the Moon and Mars. It is a "manager of managers." Such software systems are very complex; however, they are needed to make extended and complex systems humanly manageable. One of its greatest values is that the data are collected at a central location, and automatic fault management processes may be introduced, expanded, and modified as experience is gained with the network.

B. Protocol Architecture

A protocol architecture describes message formats for reporting management data and defines the managed objects. The managed objects are defined in terms of a management information base (MIB). The MIB includes a methodology for registering, identifying, and defining managed objects. There are numerous network management architectures and protocols available commercially; however, only two can be considered "standard." The first is the Simple Network Management Protocol (SNMP), developed by the Internet Activities Board (IAB) for use in the Internet, the world's largest public-access network; the second is the Common Management Information Service and Protocol (CMIS/CMIP) developed by the International Organization for Standardization (ISO) for use in Open Systems Interconnection (OSI) networks.

1. Internet Standards. The Internet has grown, especially in the last few years, as a result of the widespread availability of software and hardware supporting the Internet protocol suite. The suite includes such protocols as the Transmission Control Protocol (TCP) and the Internet Protocol (IP). The Internet Simple Network Management Protocol was demonstrated in 1988 and has been implemented by a large (and increasing) number of vendors. Details about the protocol are described in Internet documents known as Requests for Comments (RFCs) that are referenced in the following section. At the present time, virtually every new commercial router and bridge incorporates an agent based on the SNMP protocol. Many computer companies, including Sun and Hewlett Packard, have included an SNMP agent in their operating system software, which, in addition to managing communications-related objects, also enables management of computer resources, such as disk storage availability and central processing unit utilization.

a. SNMP. The SNMP architectural model involves a *network management station* and a collection of *network elements* and remote *probes* (Fig. 4). The network management station monitors and controls network ele-

ments. Network elements are devices, such as workstations, routers and the like, which have software agents. Agents perform the functions requested by a remote network management station, and act solely on the elements they reside in. Probes passively monitor the network media and measure such characteristics as total throughput and efficiency. Management stations and network elements communicate using the SNMP message protocol.

Probes are a relatively new extension to the SNMP agent architecture. Currently these probes are available for Ethernets, and they passively collect statistics and historical information from the network.

Monitoring of the network state at any significant level of detail is normally accomplished by polling for appropriate information by the network management station. A spontaneous message (called a "Trap") is used by the agent to notify a network manager of abnormal conditions.

b. MIB. In the SNMP view, the objects to be managed are identified in a management information base. The MIB is a *virtual* store—that is, a concept that identifies all the objects that need to be managed in the network along with their parameters. In its actual implementation, the values of the MIB objects may be locally stored in the element and then reported upon request to the network manager and stored in its database. The organization and structure of the MIB is described in RFC 1155 [2].

The MIB defined by the Internet community has over 100 formal objects, called the *common* MIB [3]. The MIB has recently been extended (MIB II); the examples drawn in this article relate to the common MIB. A typical device may also have an additional 100 to 200 objects that have been defined in private or experimental MIB space. This extensibility could, for example, be applied to manage elements processing Consultive Committee for Space Data Systems (CCSDS) protocols.

The common MIB is organized into eight object groups (Table 2). While details of these groups are fully described in RFC 1213 [4], Table 3 presents an excerpt from the Interfaces Group to illustrate the concept. Many of the objects are similar to what would be managed in the Earth domain of the primary network.

c. SNMP messaging. The writing (setting) and reading (getting) of variables in an agent is accomplished through the use of the SNMP message. This protocol is described in RFC 1157 [3]. SNMP models all management agent functions as alterations or inspections of variables.

The SNMP message is contained in an SNMP protocol data unit (PDU). There are five types of SNMP PDUs:

GetRequest, GetResponse, GetNextRequest, SetRequest, and Trap, and they are described in Table 4. The User Datagram Protocol is used to deliver the SNMP PDUs.

d. Remote monitoring MIB. The remote monitoring (RMON) MIB is an extension of the SNMP MIB that applies to probes. The RMON MIB for Ethernet networks has been the first to be standardized [5]. It is intended that future versions of the RFC will define extensions for other network types.

2. International Standards. International standards have also been developed by ISO for the management of OSI-based networks. The OSI management framework is more elaborate than SNMP, but is similarly designed to control, coordinate, and monitor network resources. Although the development of the OSI network management architecture is close to completion, there are very few OSI agents implemented at the present time.

The United States Government has mandated the use of selected OSI protocols for Federal information systems in the Government OSI Profile [6]. A transition to these protocols is planned for NASA administrative systems [7]. Another OSI-related profile, the Government Network Management Profile [8], has been proposed for Federal systems that specifically focuses on network management. The impact of these profiles on stimulating development of OSI network management capability remains to be seen.

OSI management protocols are generic and may be used in any OSI command and control environment (not just network management). The management structure is defined in ISO 7498-4 [9]. In the OSI architecture, when an application process, such as an agent or manager, needs to exchange information and commands with another application process, it makes use of software known as the Common Management Information Service Element (CMISE). The CMIS standard [10] defines the service that the CMISE provides and CMIP [11] defines the protocols that it uses.

CMIP and CMIS standards were published in May 1990. They provide a flexible framework for the control and exchange of management information. Together, CMIS and CMIP define the bulk of the OSI network management protocol. Revised versions of CMIS and CMIP (CMIS/CMIP Version 2) were published in January 1992. With this update, CMIS and CMIP are expected to remain stable for a number of years.

a. OSI management structure. The requirements for OSI network management are grouped into the following five major functional areas: (1) fault management, (2)

accounting management, (3) configuration management, (4) performance management, and (5) security management. While the scope is very impressive, a commercial product that implements these capabilities is not currently available.

In brief, fault management encompasses fault detection, isolation, and correction of abnormal operation of the OSI environment. Accounting management enables charges to be established for the use of resources in the OSI environment, and for costs to be identified. Configuration management includes functions to change the configuration of the system, set the parameters that control the routine operation of the system, and initialize and close down managed objects.

Performance management includes functions to determine system performance under natural and artificial conditions, to gather statistical information, and to maintain and examine historical logs. The purpose of security management is to support the application of security policies.

b. CMIS. The major difference between SNMP and OSI architectures is that while SNMP assumes an agent-manager relationship, OSI does not assume any management hierarchy. In fact, peers may communicate among themselves using CMIS/CMIP. Any necessary hierarchy is imposed by the management architects. This should simplify the development of an integrated network management (manager-to-manager) system.

CMIS management operations include a number of services that are summarized in Table 5. Note that "gets" and "sets" are similar in function to their SNMP counterparts.

c. CMIP. CMIP is a general-purpose protocol that supports the services defined by CMIS. CMIP, in turn, requires specific support from several relatively hidden OSI protocols. For example, the association services, M-INITIALIZE, M-TERMINATE, and M-ABORT, are supported by CMIP by invoking the Association Control Service Element defined in ISO 8649. The notification and operation services require use of the Remote Operations Service Element defined in ISO 9072-1. These complexities are characteristic of OSI implementations.

3. Summary of Data Architectures. The key words that summarize SNMP are "short-term" and "simple." The protocol was implemented within a few months. Today the protocol is almost universally implemented in network components, such as routers, and in certain computers. However additional extensions are being developed

in a newer version, SNMP Version 2, such as: (1) an authentication scheme to ensure security and filter out messages that may cause catastrophic errors and (2) the means to support an integrated network management system at a higher level.

As has been mentioned earlier, few, if any, OSI network management agents have been implemented in commercial equipment. However, there is strong bureaucratic interest in its adoption. The Federal Government promotes it and targets integrated network management systems as the key beneficiary.

IBM announced in March 1992 [12] that it plans to use CMIP to send information from its Advanced Peer-to-Peer Networking Network Nodes to Netview (the IBM network management system). This application of CMIP does not use OSI transport protocols; it runs over a traditional IBM Systems Network Architecture protocol stack and is thus a "hybrid" application.

The Internet community has also proposed a hybrid CMIP implementation. The implementation, called CMIP over TCP/IP or CMOT, has been implemented by a few vendors in their network manager software (not in any agents) and runs over Internet transport protocols [13].

As systems grow to include large numbers of monitored objects and subsystems, any standard evolvable management architecture remains a challenge. Fault management will be easier to scale because it usually operates on an exception basis. Performance and configuration management are more likely to initiate periodic reporting and create more traffic as the number of objects increases. The Space Management Network needs prototyping and analysis in this area.

The present application software interfaces have limitations. There is no sense of time other than "now." This makes it impossible to directly issue queries for historical information, or to issue scheduled command requests; these queries must be made through user-developed application software.

C. Connectivity

Connectivity between the element managers and the network management system may be over the primary data network or over dedicated links. An advantage of dedicated out-of-band circuits is that when the primary media becomes unavailable because of congestion or malfunctioning communications equipment, management systems can still determine and resolve the problem.

Path diversity is a major consideration in providing a wide range of options for the design of robust networks.

Both SNMP and CMIP/CMIS have control messages that can be used to select alternate paths if the initial path is blocked or highly congested.

D. Human Interfaces

The human interface provides a location for management control of the primary network. It should include a standard presentation format at the user and the systems level. It should also include the availability of artificial intelligence to assist the managers.

A variety of technologies are expected to provide improved methods for allowing users to interface to computer systems. These interface technologies focus on improvement of the amount of information that the user can perceive from a given interface configuration. Graphic visualization and interactive displays are two particularly helpful technologies. Graphic visualization may be used to represent multidimensional data on computer graphics displays in images and in a form that allows people to perceive, amplify, and interpret the data. Animated models may be used for this interface.

Interaction between the user and the data will facilitate fault resolution. Transformations and algorithms may be used to explore the effects on the data.

Hypermedia software technology enables a user to retrieve data in various formats in one or more display windows. The formats may include text, graphics, animation, digital audio, and video. Hypermedia lends itself to browsing and searching knowledge bases.

Artificial intelligence (AI) techniques that help translate raw data into knowledge may be applied—such as the technique of context sensitivity to filter data using dynamic thresholding. Also, AI techniques may be used to abstract information and present summaries to the user. AI can provide techniques for knowing what state the entire system is in and how ongoing activities are expected to affect that state.

IV. Space Management Architecture

A. Management Model

A suggested architecture to support the First Lunar Outpost is illustrated in Fig. 5. Figure 5(a) illustrates the primary network and Fig. 5(b) illustrates the SMN, pointing out the operations centers. The SMN hierarchy has four, rather than three, tiers (Fig. 5(c)). At the highest level is an Integrated Operations Center (IOC) for the highest level overview of the status of the extended network. At the next lower level are two major operations

centers: the first is an Earth Operations Center (EOC), to support the Earth and Space domains, and the second is a Moon Operations Center (MOC), an on-site facility to support the Moon domain. At a later time, a Mars Operations Center will be required. At the next lower level are facility managers that provide network management for a limited number of large facilities, such as the DSN. At the lowest level are numerous element managers; these managers oversee individual network components, such as bridges, routers, and computers.

1. Integrated Network Manager. An integrated network manager will provide high-level coordination and security management. It will primarily communicate with the EOC and MOC and sometimes coordinate activities. The IOC may serve as an alternate EOC in the event of an emergency. This situation must be jointly reported for action. The action may be to jointly modify the forward error correction coding algorithm to increase the level of coding.

2. Operations Managers. Operations management systems will be required at the EOC and MOC where global fault, configuration, performance, and security issues must be reconciled. A possible role could be resolving an event, such as a severed fiber trunk that causes sudden communications outage between the DSN and Mission Control. Facility managers will report the problem to the EOC for action and the EOC may automatically reroute the signals over a diverse path.

3. Facility Managers. The facility managers oversee all the elements associated with their facilities. At this layer there is substantial fault tolerance—usually through the ability to manually and automatically reconfigure active elements. Potential facilities include the lunar main base, DSN, Mission Operations, Science Operations, and Network Operations. A possible role could be detection of a security threat, i.e., a persistent hacker. Such a problem would normally be elevated to the IOC for action.

4. Element Managers. The proposed SMN architecture has many element managers at each of the facilities. The element managers normally communicate with the following: (1) a facility manager to coordinate a response to local events, such as reboots and power outages; (2) a local human interface for local maintenance; and (3) a local database to store MIB parameters. The element should have fault tolerance through its own internal design.

B. Data Architecture

The primary architecture requirement is that the element managers incorporate the same standard management agents. The current popular standard is, of course,

SNMP. Assuming that it will evolve into a more secure protocol with peer-to-peer communication capability, it is a leading candidate to prototype the higher management layers for the Space Management Network. A simple transition can be made later to CMIP, if such a transition is advantageous.

Network management can be implemented in the near term by specifying SNMP in all Earth-domain network elements—especially the DSN ground systems. The MIB can also be extended to include OSI and CCSDS protocol performance.

C. Connectivity

A major issue in the design of the SMN is the design of the physical network. The telecommunications industry has been migrating toward out-of-band or “common channel” signaling, a technique of putting management data into its own data channel separate from the primary data. In the case of LANs, this would be a separate LAN; in the case of WANs, it would be either a diverse circuit or a dedicated radio channel. In the space domain, the management data path may be a dedicated radio channel.

Another issue is the use of redundant communication paths for management data. A leased circuit may be identified for nominal conditions in the Earth domain. As a backup, an alternate path, such as a dial-up circuit, may be used to retrieve management data.

The bandwidth required for management data will normally be relatively low. To minimize cost, bandwidth-on-demand is an important technology for exploration missions for both the primary network and the management network.

D. Human Interfaces

Graphic visualization should be used to present multidimensional data in images and in a form that allows people to perceive, amplify, and interpret the data. Of particular concern is the configuration and status of the primary network. Animated models and interactive displays may also be used.

Hypermedia software technology should enable a user to retrieve data in various formats in one or more display windows. The formats may include text, graphics, animation, digital audio, and video.

Artificial intelligence techniques should also be applied in each of the operations centers to assist operations personnel to recognize and diagnose problems and develop alternative solutions.

V. Test-Bed Implementation

A prototype network management system for the DSN was configured in the DSN 2000 Information Systems Engineering Laboratory (Fig. 6). The management model was a simple two-tier hierarchy with a commercial network manager (SunNet Manager Version 1.2) in a dedicated Sun SparcStation IPX computer and element managers in each network element (routers and computers). The data architecture was based on SNMP, and connectivity of the network manager to the elements was through two Ethernets joined by a serial circuit through commercial routers.

SunNet has a graphical user interface based on OpenLook/X Windows. SunNet maintains a database of the network elements and it can be configured to periodically poll the elements for crucial information, such as the state of the interfaces.

The information that is obtained by polling can be examined with a textual browser or an elementary three-dimensional graphing package. The element database and polling capabilities of SunNet provide a low-level management station capability that can be extended by other vendors or by the user with additional effort.

The SunNet Manager software package included additional SNMP agents that could be installed on any Sun SparcStation. In addition to these agents, the laboratory had access to a public-domain SNMP agent, available from Carnegie Mellon University (CMU), which was useable (with minor modifications) on any Unix workstation. The CMU agent included source code, which allowed extension of the agent to monitor private variables on the Sun workstations. SunNet Manager can be configured to recognize and manage these extra variables.

Using SunNet Manager as a manager of the two-tier prototype network revealed several limitations in its use for Facility and Operations Center management; these limits may be addressed by additional application software or future SunNet Manager upgrades.

A relational database with number manipulation abilities would have been useful in analyzing the average traffic flowing over an interface and helpful in characterizing its behavior. A statistics program would have helped in analyzing data and determining what a "reasonable" number of errors on an interface would be. Ideally, a management station should incorporate an expert system with knowledge about the network's configuration, so that it can anticipate problems, suggest solutions, or automate this process. SunNet did not have these capabilities.

Another benefit that a more advanced network management station should offer is "intelligent" polling, which

can reduce the bandwidth required for network management. Under normal conditions, the network manager would request only a small number of management variables. If something anomalous were to be detected, the station could check additional variables to determine if a real problem exists.

There were several conclusions:

- (1) Inexpensive commercial software exists today for doing network management on Internet-standard networks. The managers are relatively inexpensive, ranging from \$1,000 to \$10,000, depending on the features and degree to which the software is bundled with a larger computer purchase. SunNet Manager was only \$1,000. The element agents are inexpensive or come free; they are usually included in most routers and many computers. The SNMP agent software is also publicly available, and may be used with older Unix and MS-DOS workstations.
- (2) Limited bandwidth over certain links, such as the DSN WAN circuits, require that the volume of management traffic be minimized, and this was possible to a certain extent with SunNet Manager by changing the polling rate. In the prototype, the agents were polled every 8 sec. A four-level hierarchical management structure, such as that proposed in this article, would tend to minimize traffic even more because each network manager would send only summary data to higher level managers.
- (3) The element agents can be modified to support additional MIB objects, such as those that would be DSN specific. These potentially include CCSDS protocol objects. SunNet Manager was highly configurable and easily adopted the new MIB objects.
- (4) The current SNMP standard does not explicitly include manager-to-manager communications; however, it would be possible for a user to develop this capability within the existing standard. SNMP Version 2, currently in the standards process, includes this ability.
- (5) It is possible to send management information over alternate circuits under the control of SNMP. In the laboratory, the computers and routers had diverse connectivity, which allowed management information to be transmitted over alternate Ethernets even if one specific network were down. This is an important design consideration for the Space Management Network. A robust, reliable network must have such alternate connections available. With a network management system, failures can be detected and data can be transmitted over alternate routes either automatically or with operator consent.

VI. Conclusions

A network management architecture was proposed for future exploration missions. The architecture has four layers of management: (1) element management for local monitor and control of communications nodes, (2) facility management for managing major sites, (3) two global operations centers to oversee Earth and Moon operations, and (4) an Integrated Operations Center to oversee end-to-end (e.g., lunar instrument-to-principal researcher) communications.

The most widely available standard in commercial networking hardware is the SNMP Internet standard. The standard is being upgraded by the IAB to include much of the functionality of its OSI counterpart, CMIP. Substantial cost savings will result from using the standard in the Earth domain that has the widest implementation and lowest cost, and this appears to be SNMP at the present time.

In the context of the end-to-end data system, the DSN is one of several facilities in the architecture. Although the DSN is one of the most important elements, it currently lacks a standard network management capability.

The following relevant issues remain:

- (1) DSN Network Management. The DSN facility has strategic importance as the primary Earth- and space-domain interface. A DSN management archi-

itecture, when implemented, could be used as a reference implementation for other facilities in the SMN.

- (2) Messaging Standard. Manager-to-manager communications need to be prototyped with a messaging technique, such as SNMP 2 or CMIP.
- (3) MIB Development. CCSDS standards are to be applied to most future space missions. In order to manage network processing and distribution of CCSDS packets, a CCSDS MIB is required. Based on the experience of the Internet, the CCSDS MIB could be developed and tested in a CCSDS testbed, and then be permitted to evolve based on experience. Standardization of the MIB with the CCSDS would be a necessary step.
- (4) First Lunar Outpost. A space management network for a specific proposed mission, such as the First Lunar Outpost, needs to be baselined in which specific roles for element, facility, operations, and integrated management are defined. The requirements for fault, performance, configuration, and security management need to be defined.
- (5) Automated Functions. Automated functions need to be identified for selected remote activities, and then prototyped. This capability would demonstrate configuration management responsibilities of the SMN. An example could be to simulate alternate paths between major nodes (path diversity) and include automation as a means to select an optimal path.

References

- [1] *America at the Threshold: Report of the Synthesis Group on America's Space Exploration Initiative*, Arlington, Virginia: The Synthesis Group, May 1991.
- [2] *Structure and Identification of Management Information for TCP/IP-Based Internets*, RFC 1155, Menlo Park, California: SCI International, Network Information Center, May 1990.
- [3] *A Simple Network Management Protocol (SNMP)*, RFC 1157, Menlo Park, California: SCI International, Network Information Center, May 1990.
- [4] *Management Information Base for Network Management of TCP/IP-Based Internets: MIB II*, RFC 1213, Menlo Park, California: SCI International, Network Information Center, March 1991.
- [5] *Remote Network Monitoring Management Information Base*, RFC 1271, Menlo Park, California: SCI International, Network Information Center, November 1991.
- [6] *U.S. Government Open Systems Interconnection Profile (GOSIP), Version 2.0*, Washington, D.C.: U.S. Department of Commerce, National Institute of Standards and Technology, October 1990.

- [7] *NASA Management Plan for Government Open Systems Interconnection Profile (GOSIP) Implementation*, Washington D.C.: National Aeronautics and Space Administration, OSI Management Steering Group, February 1992.
- [8] *Proposed Government Network Management Profile (GNMP), Version 1.0*, Washington, D.C.: U.S. Department of Commerce, National Institute of Standards and Technology, March 8, 1991.
- [9] *Information Processing Systems—Open Systems Interconnection—Basic Reference Model—Part 4: Management Framework*, ISO 7498-4, New York: American National Standards Institute, November 15, 1989.
- [10] *Information Technology—Open Systems Interconnection—Common Management Information Service (CMIS) Definition*, ISO/IEC 9595, New York: American National Standards Institute, April 1, 1991.
- [11] *Information Technology—Open Systems Interconnection—Common Management Information Protocol (CMIP)—Part 1: Specification—Part 2: Protocol Implementation Conformance Statement (PICS) Proforma*, ISO/IEC DIS 9596, New York: American National Standards Institute, December 22, 1988.
- [12] S. Gibson, "Products to Get APPN; Mgnt Will Use CMIP," *Communication Week*, no. 396, p. 1, March 30, 1992.
- [13] *The Common Information Services and Protocol over TCP/IP (CMOT)*, RFC 1095, Menlo Park, California: SCI International, Network Information Center, April 1990.

Table 1. Domains and potential subnetworks.

Domains	Types of subnets
Earth	Earth surface-to-surface via wire/fiber Earth surface-to-surface via radio Earth surface-to-surface via satellite
Space (Moon)	Earth surface-to-Moon surface (main base) Earth surface-to-Moon surface (remote)
Space (Mars)	Earth surface-to-Mars surface (main base) Earth surface-to-Mars surface (remote) Earth surface to Mars relay satellite
Moon	Moon surface-to-surface via wire/fiber (local) Moon surface-to-surface via radio (remote)
Mars	Mars surface-to-surface via wire/fiber (local) Mars surface-to-surface via radio (remote) Mars surface-to-relay satellite Mars relay satellite-to-relay satellite

Table 2. SNMP managed object groups.

Object group	Description
System	Objects that describe high-level characteristics of this network element
Interfaces	Objects associated with the network interfaces to which this network element can communicate with IP datagrams
Address translation	Translation table for converting an IP address into a subnetwork-specific (physical) address
Internet Protocol (IP)	Subgroup of objects associated with IP
Internet Control Message Protocol (ICMP)	Subgroup of objects associated with ICMP
Transmission Control Protocol (TCP)	Subgroup of objects associated with TCP
User Datagram Protocol (UDP)	Subgroup of objects associated with UDP
Exterior Gateway Protocol (EGP)	Subgroup of objects associated with EGP

Table 3. Representative objects in the SNMP Interfaces Group.

Object	Definition	Access
ifNumber	The number of network interfaces (regardless of their current state) present on this system.	Read-only
ifIndex	A unique value for each interface. Its value ranges between 1 and the value of ifNumber. The value for each interface must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.	Read-only
ifDescr	A textual string containing information about the interface. This string should include the name of the manufacturer, the product name and the version of the hardware interface.	Read-only
ifType	The type of interface, distinguished according to the physical/link protocol(s) immediately "below" the network layer in the protocol stack.	Read-only
ifMtu	The size of the largest datagram which can be sent/received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.	Read-only
ifSpeed	An estimate of the interface's current bandwidth in bits per second. For interfaces which do not vary in bandwidth or for those where no accurate estimation can be made, this object should contain the nominal bandwidth.	Read-only
ifAdminStatus	The desired state of the interface.	Read-Write
ifOperStatus	The current operational state of the interface.	Read-only

Table 4. SNMP PDUs.

SNMP PDUs	Description
GetRequest	At the initiation of this PDU by the sender, this PDU requests the current status of objects from the destination system.
GetResponse	In response to the GetRequest-PDU, the destination system returns the name and value of each object requested using this PDU.
GetNextRequest	This PDU is used to simplify the retrieval of successive variables in the MIB that are ordered in the form of a table (such as routing table data).
SetRequest	Upon receipt of this PDU, for each object named in the PDU, the corresponding value is assigned to the variable. The receiving station returns a GetResponse-PDU of identical form as an acknowledgment.
Trap	Upon receipt of the Trap-PDU, the data contents are passed to the application-level software for appropriate processing. Several generic traps include: <ol style="list-style-type: none"> (1) ColdStart Trap—the sending entity is reinitializing itself; its implementation may be altered. (2) WarmStart Trap—the sending entity is reinitializing itself; its implementation will not be altered. (3) LinkDown Trap—the sending entity recognizes a failure in one of the communication links represented in its configuration. (4) LinkUp Trap—the sending entity recognizes that one of the communication links represented in its configuration has come up.

Table 5. Management Operation Services.

Service primitive	Description
M-GET	Request the retrieval of management information from a peer CMISE-service-user. The service may be requested only in a confirmed mode, and a reply is expected.
M-SET	Request the modification of management information by a peer CMISE-service-user. The service may be requested in a confirmed or a non-confirmed mode. In the confirmed mode a reply is requested.
M-ACTION	Request a peer CMISE-service-user to perform an action. The service may be requested in a confirmed or a non-confirmed mode. In the confirmed mode, a reply is expected.
M-CREATE	Request a peer CMISE-service-user to create another instance of a managed object. The service may be requested only in a confirmed mode, and a reply is expected.
M-DELETE	Request a peer CMISE-service-user to delete an instance of a managed object. The service may be requested only in a confirmed mode, and a reply is expected.

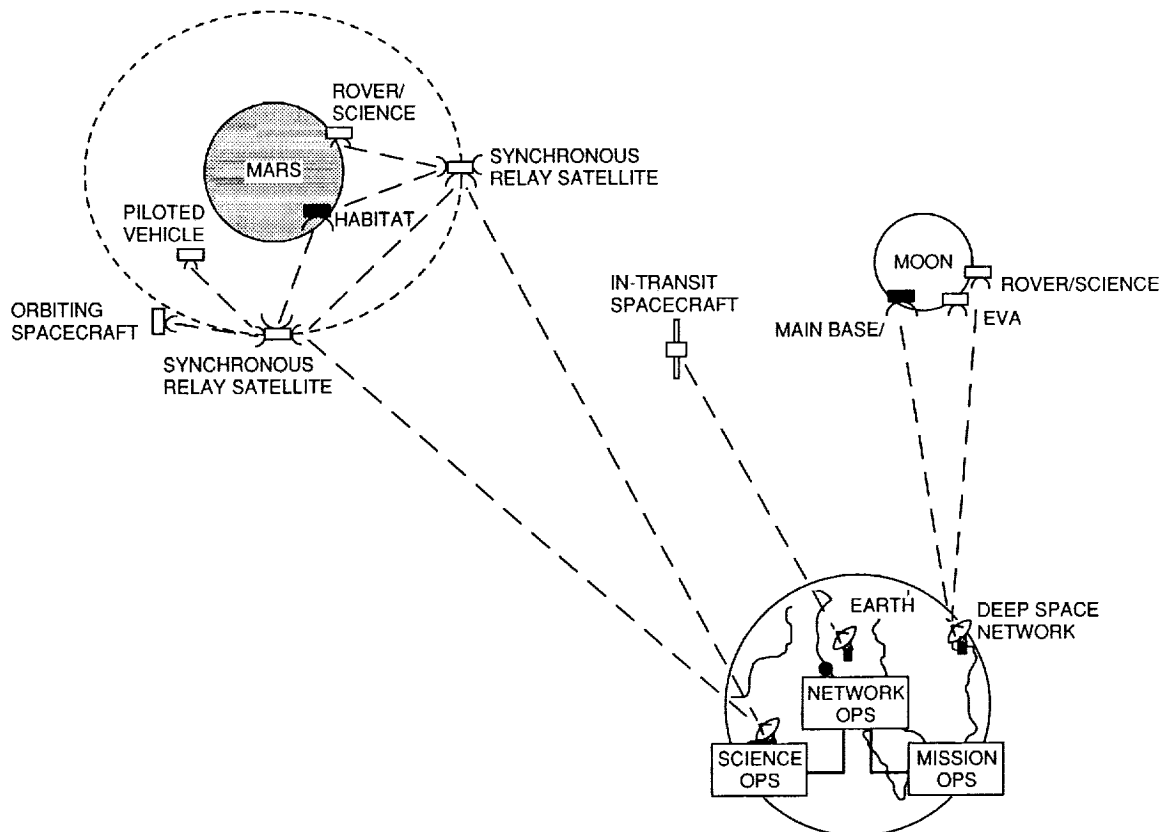


Fig. 1. Space Exploration Initiative communications architecture.

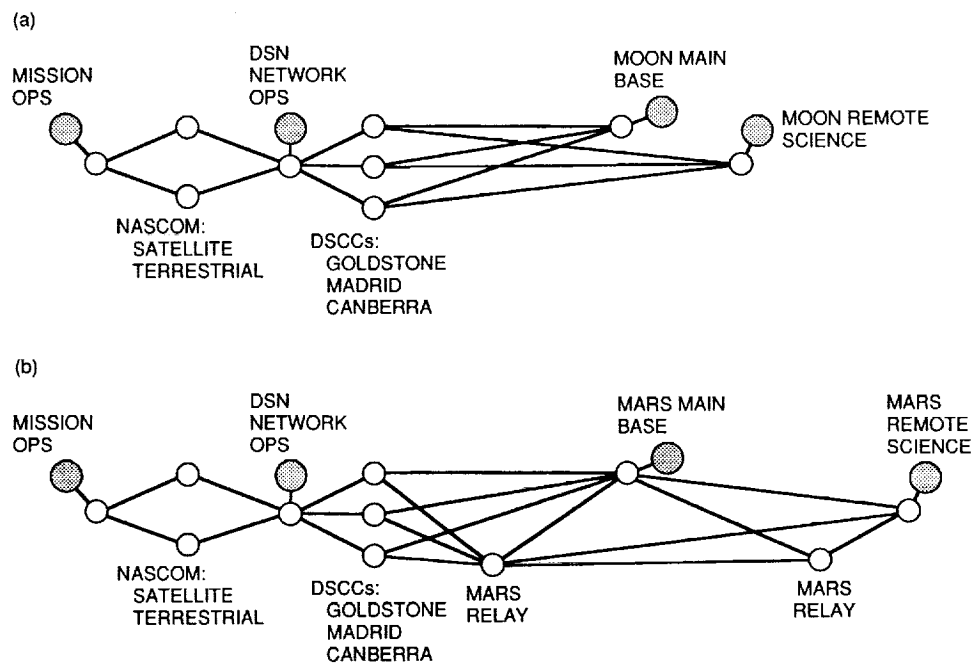


Fig. 2. Simplified schematic of primary network: (a) Moon primary space network and (b) Mars primary space network.

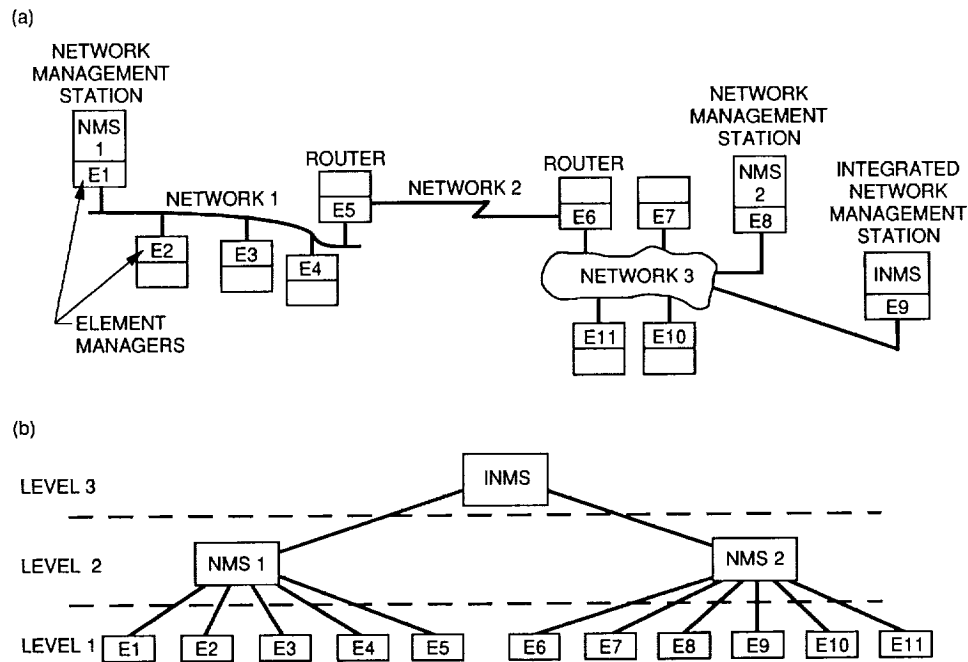


Fig. 3. Network management hierarchy: (a) physical network and (b) management hierarchy.




Fig. 4. SNMP architectural model.

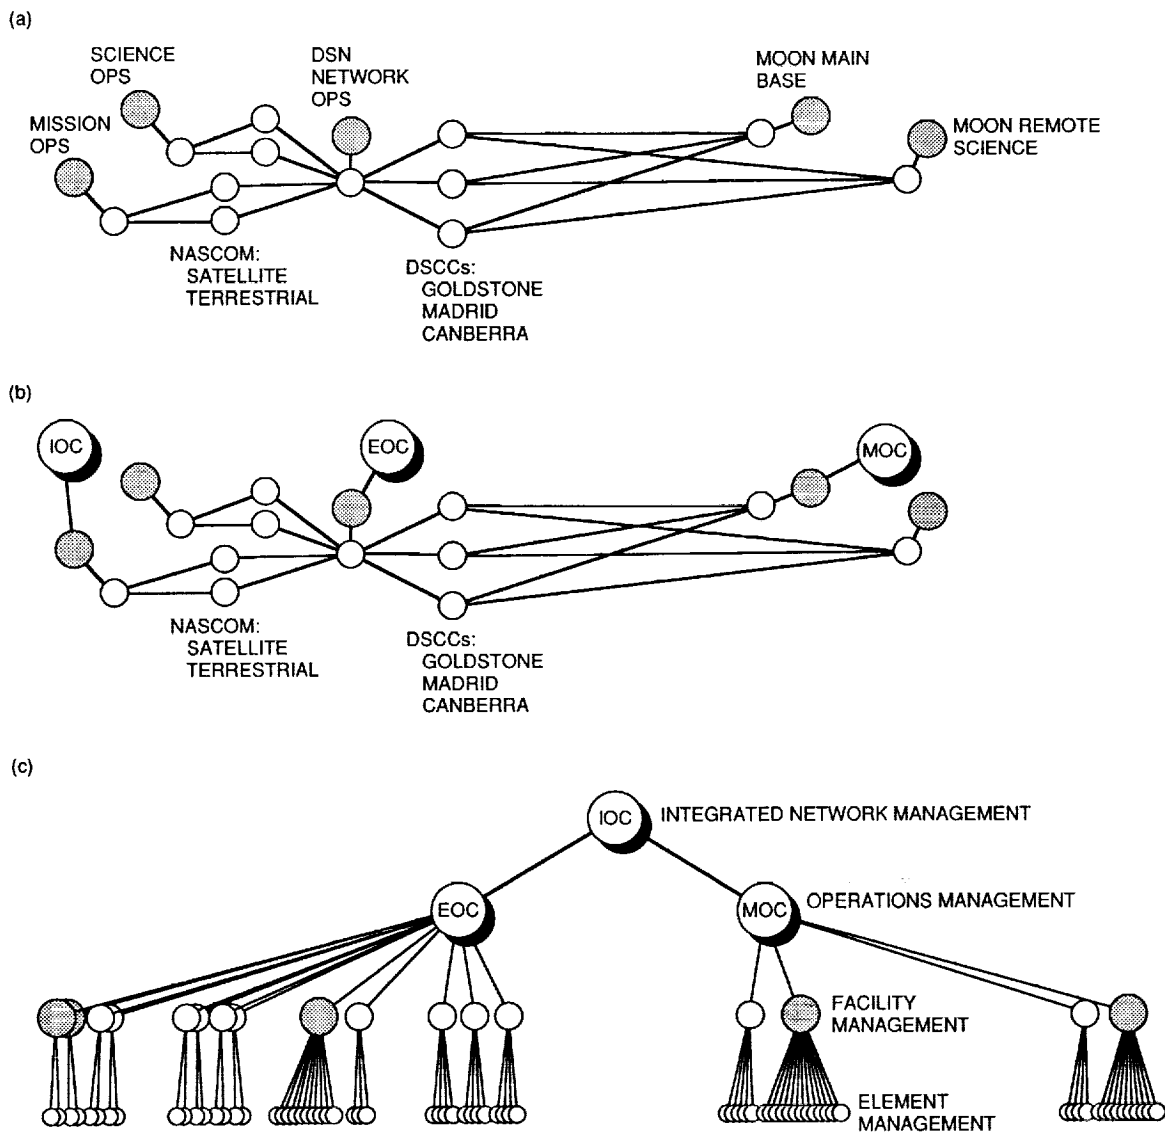


Fig. 5. Network management architecture for the Moon: (a) Moon primary space network; (b) space management network; (c) hierarchical connectivity.

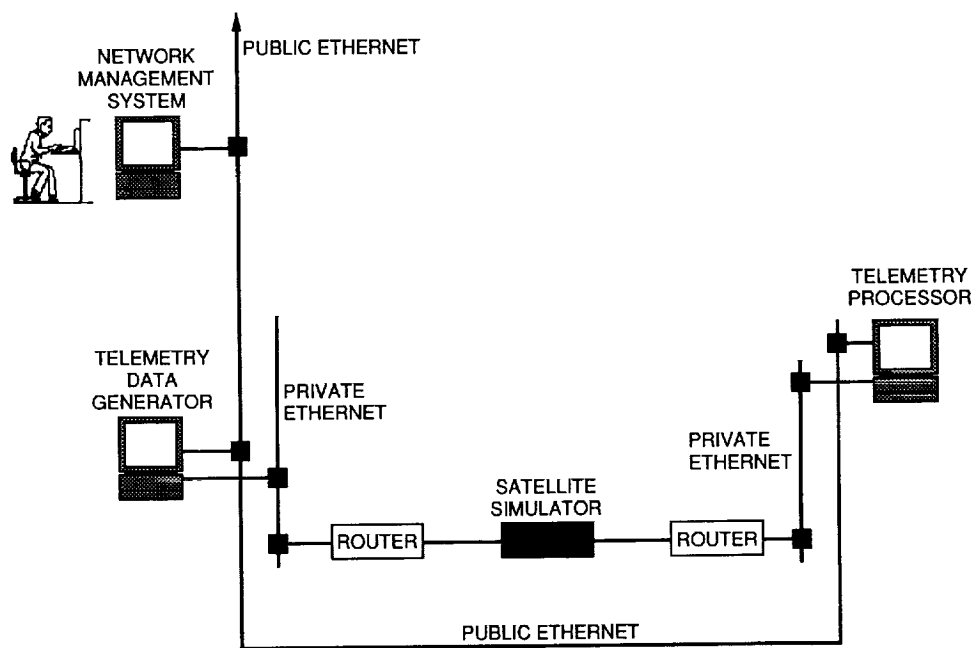


Fig. 6. DSN 2000 Information Systems Engineering Laboratory.